Thunderstorm HypOthesis Reasoner — Final Report

MTR 94B0000071
February 1994

Alice M. Mulvehill

MITRE
Bedford, Massachusetts

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	February 1994	

4. TITLE AND SUBTITLE
Thunderstorm HypOthesis Reasoner - Final Report

5. FUNDING NUMBERS
F19628-94-C-0001

6. AUTHOR(S)
Alice M. Mulvehill

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730-1420

8. PERFORMING ORGANIZATION REPORT NUMBER
MTR 94B0000071

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
NASA

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

THOR is a knowledge-based system which incorporates techniques from signal processing, pattern recognition and artificial intelligence (AI) in order to determine the boundary of small thunderstorms which develop and dissipate over the area encompassed by the Kennedy Space Center and the Cape Canaveral Air Force Station. THOR interprets electric field mill data (derived from a network of electric field mills) by using heuristics and algorithms about thunderstorms that have been obtained from several domain specialists. THOR generates two forms of output: contour plots which visually describe the electric field activity over the network and a verbal interpretation of the activity. THOR uses signal processing and pattern recognition to detect signatures associated with noise or thunderstorm behavior in a near real time fashion from over 31 electrical field mills. THOR's AI component generates hypotheses identifying areas which are under a threat from storm activity, such as lightning. THOR runs on a VAX/VMS at the Kennedy Space Center. Its software is a coupling of C and Fortran programs, several signal processing packages, and an expert system development shell.

14. SUBJECT TERMS	15. NUMBER OF PAGES
thunderstorm, artificial intelligence, electric field mill data	26
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	Unlimited

Thunderstorm HypOthesis Reasoner — Final Report

MTR 94B0000071
February 1994

Alice M. Mulvehill

Accesion For		
NTIS CRA&I		☒
DTIC TAB		☐
Unannounced		☐
Justification		
By		
Distribution /		
Availability Codes		
Dist	Avail and / or Special	
A-1		

Contract Sponsor NASA
Contract No. F19628-94-C-0001
Project No. 586I
Dept. G041

MITRE

Bedford, Massachusetts

Department Approval: ______________________

Mark T. Maybury, Department Head

MITRE Project Approval: ______________________

Project Leader, 3138I

ABSTRACT

THOR is a knowledge-based system which incorporates techniques from signal processing, pattern recognition and artificial intelligence (AI) in order to determine the boundary of small thunderstorms which develop and dissipate over the area encompassed by the Kennedy Space Center and the Cape Canaveral Air Force Station. THOR interprets electric field mill data (derived from a network of electric field mills) by using heuristics and algorithms about thunderstorms that have been obtained from several domain specialists. THOR generates two forms of output: contour plots which visually describe the electric field activity over the network and a verbal interpretation of the activity. THOR uses signal processing and pattern recognition to detect signatures associated with noise or thunderstorm behavior in a near real time fashion from over 31 electrical field mills. THOR's AI component generates hypotheses identifying areas which are under a threat from storm activity, such as lightning. THOR runs on a VAX/VMS at the Kennedy Space Center. Its software is a coupling of C and Fortran programs, several signal processing packages, and an expert system development shell.

ACKNOWLEDGMENTS

Funding for this project was provided through contract number: NAS918057. The system was developed through the efforts of The MITRE Corporation, Bedford, MA; NASA at the Kennedy Space Center; and EG&G at the Kennedy Space Center. This document was written under the guidance and suggestions of the NASA domain specialists Launa Maier and James Nicholson. Special thanks to Richard Neely of EG&G for his work on the system and to Tom Gould for editing a preliminary version of this document.

TABLE OF CONTENTS

LIST OF FIGURES

SECTION 1

INTRODUCTION

THOR, which stands for "Thunderstorm HypOthesis Reasoner" is a knowledge-based approach to thunderstorm forecasting. The system has been developed to detect the increased electrification associated with the development of small, slow moving (if not stationary) summer thunderstorms which occur in the Kennedy Space Center (KSC) and Cape Canaveral Air Force Station (CCAFS) area, and to delimit the boundary of the electrification.

Figure 1 shows the KSC and CCAFS region. Each number of the map corresponds to a field mill location; thus, number 27 corresponds to field mill 27. Several work enclaves are spread throughout this region, and each can be adversely affected by increased atmospheric electricity and lightning from thunderstorms. These areas include the industrial complex, the landing site, the VAB (Vehicle Assembly Building) and OPF (Orbiter Processing Facility) area, Launch pads 39A and 39B, and the Cape launch sites.

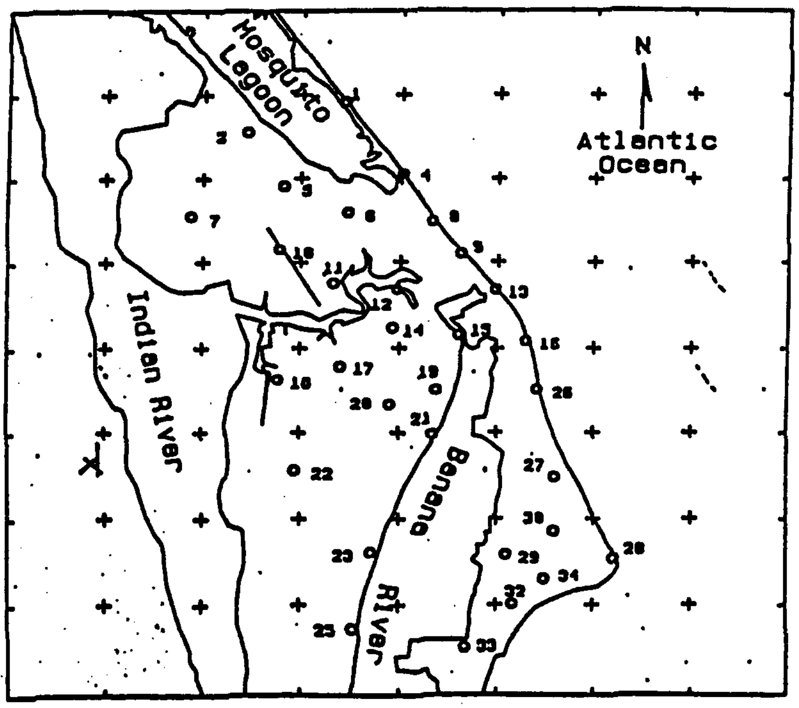

Figure 1. Map of the KSC/CCAFS Area

Lightning storms and electrified clouds, especially prevalent during the summer months, present a serious threat to the safety of personnel and equipment, along with timely and costly work disruptions. Forewarning needs vary by operation, from five to thirty minutes, in order for personnel to assure a safe condition. Complicating the problem is Central Florida's high lightning frequency. With Cape Canaveral's adjacent broad water and land areas, and plenty of moisture available in the boundary layer, the surface thermals produce upcurrents, resulting in possible daily thunderstorms.

A local storm varies in duration from thirty minutes to four hours, and the number of lightning flashes can vary from one to several thousand. Ground contact area ranges from one to hundreds of square kilometers. Ground lightning presents a serious threat to local operations. Frequently summertime thunderstorms grow in place, are relatively small, and remain stationary or migrate slowly; and the fury of a small thunderstorm often belies its size.

SECTION 2

THE ELECTRIC FIELD MILL NETWORK

NASA maintains a network of ground-based electrical field mills (Figure 1) for identifying clouds capable of producing electrical hazards to space vehicles and ground operations. The field mill sensors provide data on the electrification of clouds through the storm's lifecycle. The mills are acutely sensitive to changes in the electric field, and their geographical distribution provides full coverage of the KSC/CCAFS area.

Weather forecasters use field mill data to locate specific areas of significant ground electric field disturbances. Forecasters also use the mills to detect lightning flashes, locate them, and determine the development and decay of thunderstorms.

A thunderstorm progresses through several phases during its lifecycle, such as fairweather, increasing electrification, active and decreasing electrification. The Thunderstorm Identification Process diagram show in Figure 2 depicts a typical flow. During each phase, specific electrical events called *signatures* can be identified and used to describe and even predict thunderstorm behavior. Thunderstorm signatures include: initial zero crossing, increasing trend, lightning, decreasing trend, foul to fair zero crossing, and end-of-storm-oscillation (EOSO).

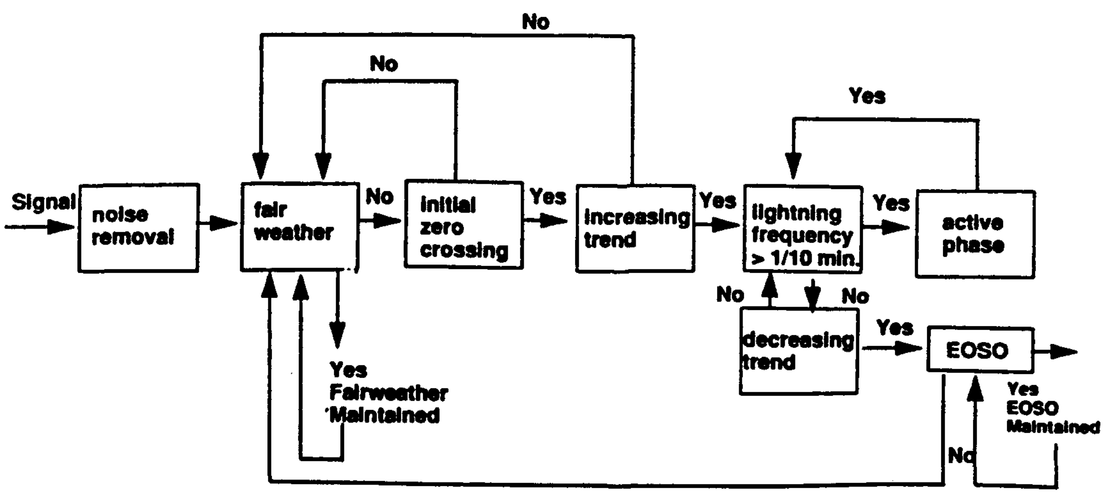

Figure 2. Thunderstorm Identification Process

Weather forecasters view depictions of data from field mills through contour plots and strip charts and through other systems such as radar to draw inferences for mesoscale weather predictions. Too much data, however, can burden the forecaster's ability to make fast predictions.

SECTION 3

THOR

THOR is a prototype system which employs a combination of signal processing, pattern recognition, and artificial intelligence (AI) techniques. THOR is designed to operate in a near real time, local forecasting environment. Its goal is to generate a hypothesis concerning the early electrical activity associated with the development of small, stationary thunderstorms and to determine the boundary of the storm. THOR currently operates on data from small, slow-moving, summer thunderstorms. THOR is envisioned to handle data from more complex thunderstorm systems.

In the past, automated weather prediction tools were developed and deployed in forecasting facilities to support data interpretation from single and multiple sources. These tools range from ones using data representation as a method for compressing data, to those that analyze data to reduce or summarize it. A few systems even attempt to analyze and interpret the data. Several of these latter tools are knowledge-based expert systems developed using AI [1, 2, 3, 4].

In developing THOR a knowledge-based systems approach was employed to capture resident field mill analysts' expertise on the development and dissipation of *stationary* thunderstorms. The development of the system was facilitated by the concentration on this one type of thunderstorm and because of the existence of:

- an array of electric field mills sufficient to reveal all or much of the thunderstorm or electrified cloud perturbations

- special concerns raised by small thunderstorms, a class of storm that develops in the locale and is nearly stationary

- domain experts willing and dedicated to supporting the design of the system.

SECTION 4

SYSTEM ARCHITECTURE AND DEVELOPMENT

THOR runs on a VAX mainframe computer under the VMS operating system. Various software packages and utility programs, written in both C and Fortran, are employed. The Data Advisor Module also runs independently on a PC platform.

Two modules make up THOR's architecture, as shown in the diagram of Figure 3. The Data Advisor is a knowledge-based module. It represents the reasoning process employed by domain specialists to analyze electric field mill data as a series of decision trees consisting of rules. Many of the knowledge bases use clustering and nearest neighbor algorithms to compare and validate data from multiple sensors. Several of the knowledge bases comprise the "Correlator" which is an adjunct module to the Data Advisor. The Correlator works by using output from the Data Advisor to correlate mill activity across the entire network.

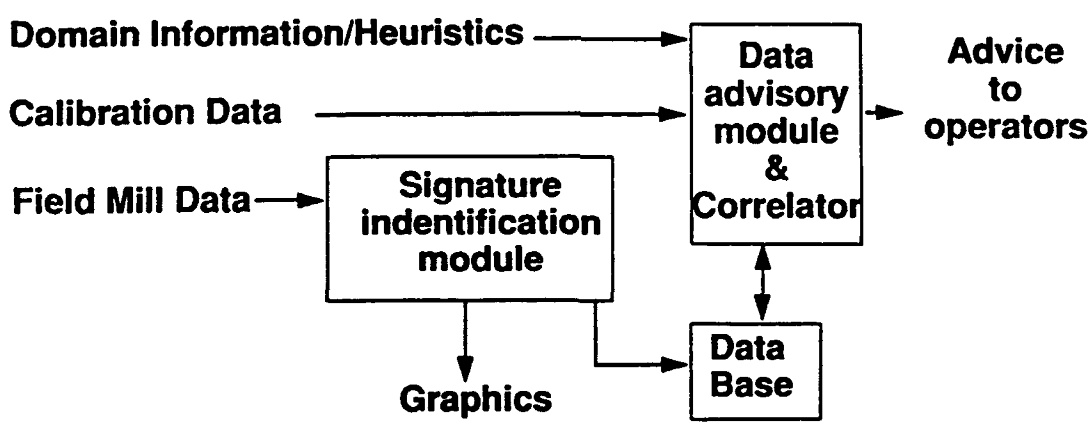

Figure 3. THOR's Architecture

The Signature Identification Module (SIM) uses signal processing algorithms and pattern recognition techniques to detect signatures from field mill data. The SIM uses two software packages: Matlab and IMSL/IDL. Matlab is employed as the environment for developing and testing signature detection algorithms. It performs interactive numeric computation, data analysis, and graphics. Matlab provides hundreds of built-in advanced functions, and two-dimensional (2-d) and three-dimensional (3-d) graphing capabilities. It features an interactive, yet fully programmable environment. Matlab also offers a family of toolboxes for specialized applications, including control system design, digital signal processing, and optimization.

IMSL/IDL (International Mathematical and Statistical Library/Interactive Data Language) is used for mathematical processing and to generate most of the graphics for the system. It provides a complete computing environment for the interactive analysis and visualization of scientific and engineering data. It integrates a powerful, array-oriented language with numerous mathematical analysis, image processing and graphical display techniques.

Events detected by the SIM are provided to the Data Advisor via a database. The database contains 65 fields. Seven of these fields contain state information about each mill in time. The seven signature fields comprise a state vector which describes the state of each mill at one minute updates. The state vector information is also used by the SIG to color the mills on the contour maps as the field mill data is being processed.

Figure 4 displays the state field name, the signature name it corresponds with, and what the values in the state vector specify. While a mill is typically **in** (1) or **out** (0) of a signature state at any given time, some signatures, like the EOSO cannot be detected for up to 20 minutes. In order for the Data Advisor to not be delayed in generating hypothesis concerning electrification, three values (0, 1, 2) are used. For example, when the *ieev* is 0, then **increasing trend** is not being evaluated or detected, when *ieev* is 1 the Data Advisor can begin to reason that the mill with *ieev* as 1 is beginning to show increased electrification although the **increasing trend** signature has not yet been detected (as indicated by a 2 in the *ieev* vector field).

field name	signature name	value
fwev	fairweather	0 = not in fairweather 1 = in fairweather
izev	initial zero crossing	0 = no zero crossing detected 1 = zero crossing detected
ieev	increasing trend	0 = no increasing trend detected 1 = checking for increasing trend 2 = increasing trend detected
liev	lightning	0 = no lightning detected 1 = lightning detected
deev	decreasing trend	0 = no decreasing trend detected 1 = checking for decreasing trend 2 = decreasing trend detected
fzev	foul to fair zero crossing	0 = no crossing detected 1 = zero crossing detected
eosoev	EOSO	0 = no EOSO detected 1 = checking for EOSO 2 = EOSO detected

Figure 4. State Vectors

While the state vector information is used by the Data Advisor to check for current state changes; the Data Advisor uses other data fields, e.g., *twoavg* (two minute average), *avgslope* (average slope) and *lsum* (lightning sum for past 10 minutes) in order to check for the occurrence of an event in the past and in order to check for persistence (as in the case of the lightning count over the past 10 minutes). This is necessary for example, to recognize that a mill crossed zero and satisfied the *increasing trend* slope in the past, which are events the Data Advisor must check for before it will begin to check for the *active* phase. This requirement to check for past events is directly related to the weak model that THOR uses to reason about thunderstorm behavior. The model simply states that the following storm phases occur in order: fairweather, increased electrification, active phase, and dissipation.

The Data Advisor generates a hypothesis about the electric field and atmosphere in the vicinity of a sensor using this data. The Correlator takes as input the hypotheses output by the Data Advisor for each sensor and generates a classification tree which describes the entire field by time and or by hypothesis (typically some thunderstorm state such as fairweather).

The Data Advisor and the Correlator both run under a tool called First Class Fusion/HT (Note: HT stands for hypertext; however hypertext is only available on the PC version of the tool). This is an inductive-based expert system development tool that uses the ID3 learning algorithm to induce a production rule (expressed as a decision tree) [5, 6,]. Examples are distinct attribute-value sets and can be viewed as the training set for developing the decision tree. ID3 analyzes data by looking for similarities within a set of examples or cases.

SECTION 5

USING THE SYSTEM

THOR has been developed to better determine the areas that could be affected by the development of a small, stationary thunderstorm. The storm of 24 September 1986 is an example of this type of storm. Data recorded from this storm was used to develop the signature detection algorithms and many of the knowledge bases of the Data Advisor. Data from a similar storm, 17 August 1984, was used to test both modules of the system. Performance evaluation results indicate that the system did as well on the 1984 storm as it did on the 1986 storm. A more detailed discussion of the performance evaluation results and some observed differences is provided later in Section 5.2.

5.1 THE STORM OF 24 SEPTEMBER 1986

The electric field mill data for the 24 September 1986 storm was corrupted with noise associated with pulse code modulation (PCM). Since this noise was very regular, it was classified as a signature and a signature detection algorithm was developed to recognize and remove it. The algorithm, RMNOISE was successfully used to remove the noise from the data so that the signatures and atmospheric behavior of the storm could be detected.

Analysis of the data from the 24 September 1986 storm revealed that the storm was preceded by several small showers in the area of mills 12, 15, and 17. As electrification developed, the storm became centered in the vicinity of mill 10. This activity can be monitored in THOR. For example, through the SIG, strip chart recordings produced for each mill display the amplitude of potential gradient versus time. The SIM places colored numeric markings on a strip chart to indicate the occurrence of the following signatures:

- the start and end of fairweather in white.
- the time of the Zero Crossing, or polarity reversal in cyan.
- the increasing electrification in yellow.
- lightning, which signifies the active phase, shown as triangles, colored red.
- the decreasing electrification is identified by magenta.
- the foul-to-fair zero crossing is marked by blue.
- the EOSO (end-of-storm oscillation) is orange.
- and the return to fairweather is again shown in white.

No two field mills respond exactly the same to electric field changes. This is due to such factors as the field mill's calibration status, environmental conditions in the vicinity of the mill such as interference from local ground cover or nearby wind towers, and the field mill's distance from the center of the storm. As such, atmospheric conditions at one mill might not be visible on another mill. For example, if one were using THOR as it operates today on the

VAX, one would notice that the electrical activity associated with the occurrence of a small shower visible on the strip chart for mill 12 is not evident on mill 10's strip chart.

The distance between a field mill and the center of the storm will affect the amplitude of the readings on the distant mill (assuming it is a good, recently calibrated mill). For example, mill 27 is substantially distant from mills 10 and 12 which are located under the center of this storm. A look at mill 27's strip chart shows that it stays primarily in the fairweather polarity and does not cross zero until it is affected by the intense lightning in the region near mill 10, occurring at 14:50 hours.

Alternative ways of viewing how the entire network of field mills reacts to changes in the electric field are provided by the system in the form of three-dimensional plots and contours plotted on the regional map. The 3-d plot shown in Figure 5 shows the electric field mill activity for all mills on 24 September 1986. The plot is displayed in a north-south orientation. The plot displays the occurrence of preceding showers, polarity reversals, the lightning associated with the stationary storm in the region near mill 10, the return to fairweather, and the later development of a small storm centered in the vicinity of mill 18 that started at 18:30 GMT.

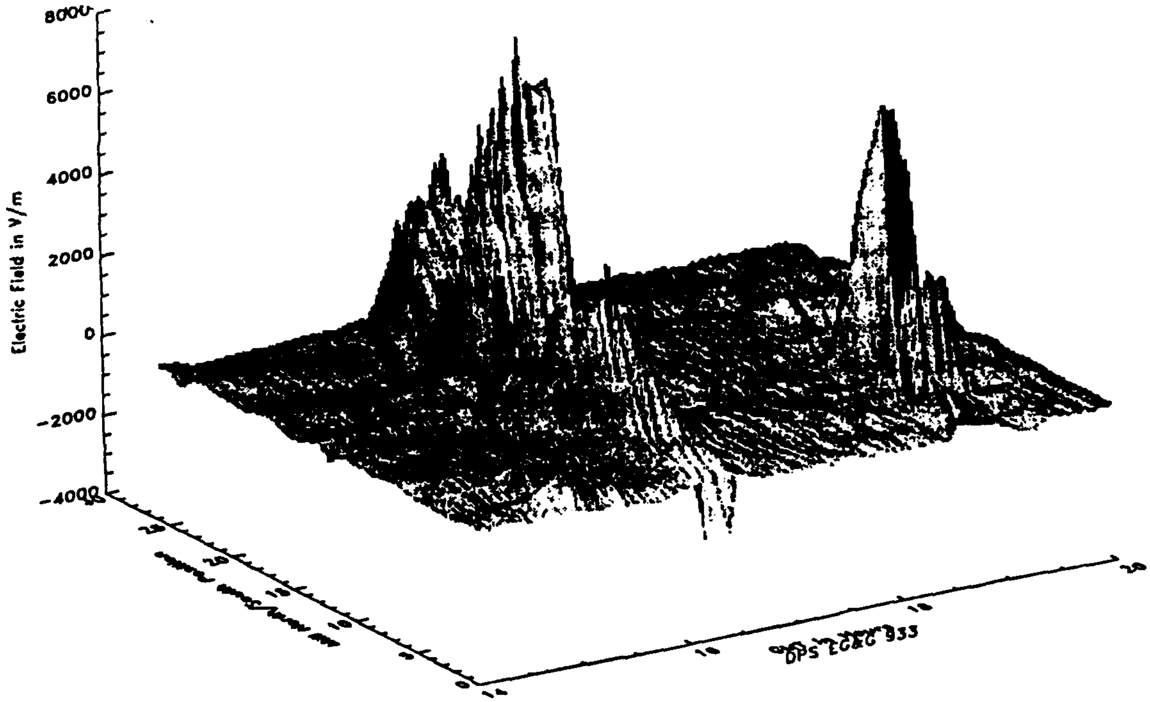

Figure 5. 3-d Plot of 24 September 1986 Storm

If one were to look at the strip chart of mill 18, one would see the activity associated with the first storm, the return to fairweather, and the more significant increase in electrification and lightning associated with the second storm which is centered in the vicinity of this mill.

While the strip charts and 3-d plots generated in THOR have been used in this project as post storm analysis tools, THOR is intended to be a real time system, processing and displaying data in a near real time fashion. Currently the system can process and display data at one minute intervals. Two map displays are employed. Each map displays the field mill locations delimited by their numbers. The color of each field mill number changes as the storm develops and decays over time. GMT is displayed in large numbers and is updated in one-minute increments. As the map displays the changes in state, indicated in color with the detection of specific signatures, the Data Advisory module simultaneously generates hypotheses about the atmosphere.

In one map display, the detection of specific signatures is communicated to the user by color, where:

- fairweather is green
- decreasing trend is yellow
- lightning is red
- active phase is orange
- decreasing trend is magenta
- and the EOSO is purple

A white triangle is located next to each field mill number and will turn red when lightning is detected by that field mill. Contours are generated to indicate the changes in the electric field and the boundary of the electrical activity. The electric field in volts/meter/minute is also displayed on the contour lines.

An alternative method of viewing the storm data is to use the Storm Electric Field Map display. Instead of showing the detection of specific signatures like *increasing trend*, the map displays colors indicative of polarity reversals. Positive electric field polarity is indicated by yellow. Negative electric field polarity is blue. The contours behave as they do on the other map, displaying the changes in the electrification intensity and boundary.

Other visualization techniques are available with THOR. These include: a colored 3-d plot, where the colors indicate electrification levels (blue is fairweather and red means high fields of lightning); and horizontal section views of a 3-d plot. For example, the horizontal section view of a 3-d plot of the 24 September 1986 storm dramatically delimits the boundaries of each of the small storms that developed, and is particularly instrumental in showing the coverage by the individual storms and the fact that neither storm showed any indication of being a moving storm.

5.1.1 An Animated Description of the Storm of 24 September 1986

One can view THOR as it processes the storm data from the 24 September 1986 storm on the VAX/VMS system at KSC, or through either the promotional videotape or the training videotape that was generated at the end of 1993. As one watches the system process the data of 24 September 1986, using, for example the map display which specifies the detection of specific signatures, a trained field mill analyst or forecaster would be able to determine what was occurring in the atmosphere from the contour shapes and voltage values. The trained analyst may or may not be able to conclude that certain contours are associated with rain without referring to the associated strip charts. Coupling the output of the Data Advisor to the contour map displays produced by the SIM alleviates this uncertainty in the trained analyst and provides qualitative interpretation to the novice analyst.

For example, the map for the storm of 24 September 1986 begins with contours indicative of off-shore activity in the southeast part of the network. The mills in this area are colored magenta to indicate that they are in a the decreasing trend, which implies that the electrification is dissipating in this region. The other mills on the map are all denoting fairweather at this time. Output from the Data Advisor confirms this behavior.

Field Mill 12 then changes to yellow, indicating a small shower. Field Mill 15 also changes color, revealing shower activity it its area. Field Mill 17 turns color and displays contours, indicating the more significant shower activity in its area. The showers decay and the storm begins to intensify in nature. While the map displays field mills 12, 15, and 17 as yellow for increased electrification, the Data Advisor hypothesizes the existence of a small shower and advised the operator to check neighboring mills and to continue to monitor the mills in the area.

As the mill numbers on the map change colors to denote the detected activity, the contour plots change to indicate the change in electrification and the boundary. The mill numbers in the vicinity of mill 10 begin to change color to show increased electrification and associated reversals in polarity or zero crossings. At this point the Data Advisor indicates a significant perturbation has occurred and checks neighbors to determine the boundary of the storm.

The white triangles then turn red to indicate lightning. Note: while THOR can detect and count lightning flashes, other forecasting systems, like LLP (Lightning Locator Processor) and LDAR (Lightning Detection Aperture Radar), currently used by forecasters at the KSC/CCAFS forecasting facility, are far more sophisticated at detecting and locating flashes. As an operational system, THOR would use data from these systems in order to indicate and evaluate lightning.

As the first storm of 24 September 1986 begins to decay, indicated by a reduction in contour plots and the change in color from red and orange to magenta and purple, the Data Advisor hypothesizes that the storm may be entering the End of Storm Oscillation, but could still

reenter the active phase. When the SIM indicates the start of EOSO, denoted by purple, the Data Advisor hypothesizes that the system is returning to fairweather.

Once back in fairweather all the mills are colored green. Fairweather is maintained until 18:30 hours. At this time, a new storm develops to the west of the network. This activity results in changes in state, denoted by color and contour plot changes centered around field mill 18. Other mills in the region also change color. As the second storm begins to develop, the Data Advisor once again generates hypotheses indicating storm activity and the map changes color to denote the changes in state. The contours close around field mill 18, showing where the storm, still west of the network, has the greatest influence. However, it is important to note that while the contour analysis is representative over KSCs field mills, it is not representative of reality in the area to the west over the Indian River, so a complete picture is lacking of the storm's influence. In other words, while the analysis delineates the boundary of the storm's influence over KSC, providing valuable information on the threat to work enclaves, THOR cannot locate the center of the storm due to lack of field mill data. The storm's center may be farther west than the center of the closed contour. Furthermore, if THOR was provided with other data about the environment, it could conclude as the domain specialist did that "because the environmental winds were extremely light on 24 September 1986 and not likely to move, this storm's threat was confined to the western part of KSC and posed no threat to pads 39A and 39B".

5.2 THE STORM OF 17 AUGUST 1984

The SIM encountered some time periods in the 17 August 1984 storm where there was corrupted data. This corrupted data, which occurred primarily within a 10 sample/second period was hand edited. Since the SIM operates on samples, it simply recorded the time shift generated by the lost data. Note: data missing on less than a one minute interval will cause no problems to the Data Advisory module, but if the data is missing for any one minute period, the system will generate a database READ error. When the system is operating unassisted in a semi-real time mode and there is a one-minute chunk of data missing, some DCL command should be called from the Data Advisor to guide it over the gap. The RMNOISE algorithm was then applied to remove any PCM noise.

Preliminary performance evaluations indicate that in most cases the **initial zero crossing** and **increasing trend** algorithms correctly detected the signature; detection of the **decreasing trend** signature was erratic; and the detection of the **EOSO** was successful about 50% of the time. A more detailed evaluation was made on the initial zero crossing values. The results in Table 1 show the values to be within one minute of crossings derived from two-minute average calculations of the storm data. Detailed evaluation of other signature identification results are not yet available.

Mill Number	SIM Zero	2-Minute Average Zero
1	n/a	n/a
2	18:23:03	18:22:42
4	n/a	n/a
5	18:35:51	18:34:46
6	18:42:09	18:41:49
7	18:36:02	18:34:46
8	18:52:44	18:52:52
9	18:56:55	18:56:54
10	18:22:20	18:22:42
12	18:36:03	18:36:47
13	18:44:27	18:42:49
14	18:22:11	18:22:42
14	18:39:58	18:39:48
15	17:57:00	17:57:34
15	18:41:11	18:41:49
16	18:41:22	18:39:48
17	18:13:40	18:13:49
18	18:04:19	18:04:36
19	18:17:10	18:17:40
20	18:14:14	18:14:40
22	18:03:00	18:03:36
23	18:04:43	18:04:36
25	18:02:14	18:02:36
27	18:32:34	18:31:45
29	19:17:37	18:17:40
30	18:32:41	18:32:46
34	18:20:28	18:21:42

Table 1. Zero Crossing Comparisons

As with the 24 September 1986 storm discussion, one can use the system to watch the storm of 17 August 1984 develop and decay through the graphical and advisory output of THOR. (Note: it is not on either of the videotapes). Figure 6 shows the 3-d plot of the activity for the 17 August 1984 storm. Notice that this storm has more lightning associated with it as evidenced by the peaks in the 3-d plot and that the activity covers a wider portion of the network for an extended amount of time.

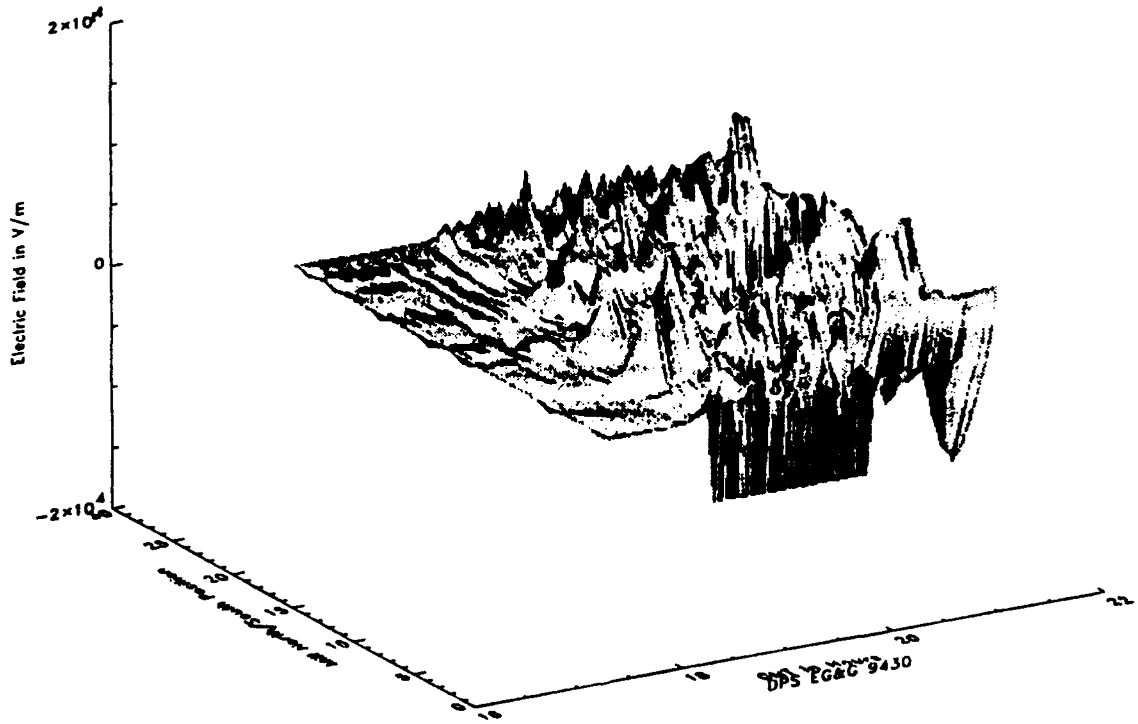

Figure 6. 3-d Plot of 17 August 1984 Storm

In terms of hypotheses and/or advice generated by the Data Advisor, evaluation of the results show that the hypotheses generated by the Data Advisor generally matched the expectations of the domain specialists and/or performance evaluators. An exception was with regard to Mill-1 and Mill-4. Visual inspection of these mills led the domain specialist to conclude that both were probably bad since both begin in a negative polarity and tend to not show much change over time. However, while the Data Advisor generated a hypothesis that Mill-1 might be a bad mill; it concluded that there was a shower of small extent in the vicinity of Mill-4. For both mills (both located along the northeast coastline) the advisor hypothesized that there was a disturbance external to the network at the time period 17:37:27. Contour plots generated by the SIM show a disturbance off the network near both Mill-1 and Mill-4 in the time period 17:37:27, the time when the Data Advisor reached this conclusion.

The interesting aspect of these hypotheses is that the system never concludes that Mill-4 is a bad mill, while the field mill domain specialist suspects that Mill-4 is a bad mill since its slope over time tends to be about zero. Currently, there is no test on the stability of the slope in the Data Advisor, therefore, the Data Advisor can not make this assessment nor arrive at this conclusion. Discussions were held to determine where such a check should be made. In the Data Advisor, this check, placed at the wrong reasoning stage could erroneously conclude that a clean mill in *fairweather* is bad because its slope would (correctly) be approximately zero for some amount of time. A suggestion is to add a data field to the database to keep track of a 10 minute slope as calculated in the SIM. Although nothing has been implemented at this time, this is an example of the type of situation that has not yet been captured by the reasoning heuristics of THOR.

SECTION 6

MODES OF OPERATION

THOR can be used as a real time system or as a post storm analysis tool. As a real time system, THOR's SIM can process the data without intervention. The Data Advisor operates through VAX/VMS DCL commands and automatically generates reports which contain all hypotheses generated during its evaluation of a mill at a particular time. As a post storm analysis tool, electric field mill strip charts, maps, and 3-d plots can be generated by the SIM upon demand. Using one of the interactive modes provided by the Data Advisor, a user can pulse through the data, requesting information for any mill by mill-name, stormdate, and Julian time (a transformed GMT). The main menu of the Data Advisor is shown in Figure 7. The **Demonstration or Testing** option is the best suited for most post storm analysis.

```
THOR
[F1=Help]                                        [Esc=Stop]
What action would you like to perform?

    Run the system (near Real Time Mode)
    Demonstration or Testing
    Interactive Mode
    Access/Modify database information
    Correlator
    Tutor
    Quit
```

Figure 7. Data Advisor Main Menu

The complete system (SIM plus the Data Advisor) does not currently have a comprehensive user interface, therefore, the display of any graphics or Data Advisor output is under the control of the user. Having control of the interface is very useful during post storm analysis, but during the real time processing of storm data, a more comprehensive interface, like that displayed in Figure 8, should be implemented. In this top-level display, the stormdate and current GMT is constantly displayed, along with one or several concurrent 2-d and 3-d graphical views of the data. In this interface, user preferences should be enabled, and any one window might take precedence over other windows at particular times. For example, during general processing, one of the 3-d contour maps would be displayed, with the stormdate and GMT information in the upper left corner of the display and hypotheses generated by the Data Advisor appearing in a window in the upper right of the screen. If the user decided to make the map larger so that it encompassed the entire screen, that would be allowed until the Data Advisor generated a hypothesis which departed from its previous hypothesis (the mill was in fairweather for the past hour and is now showing shower activity).

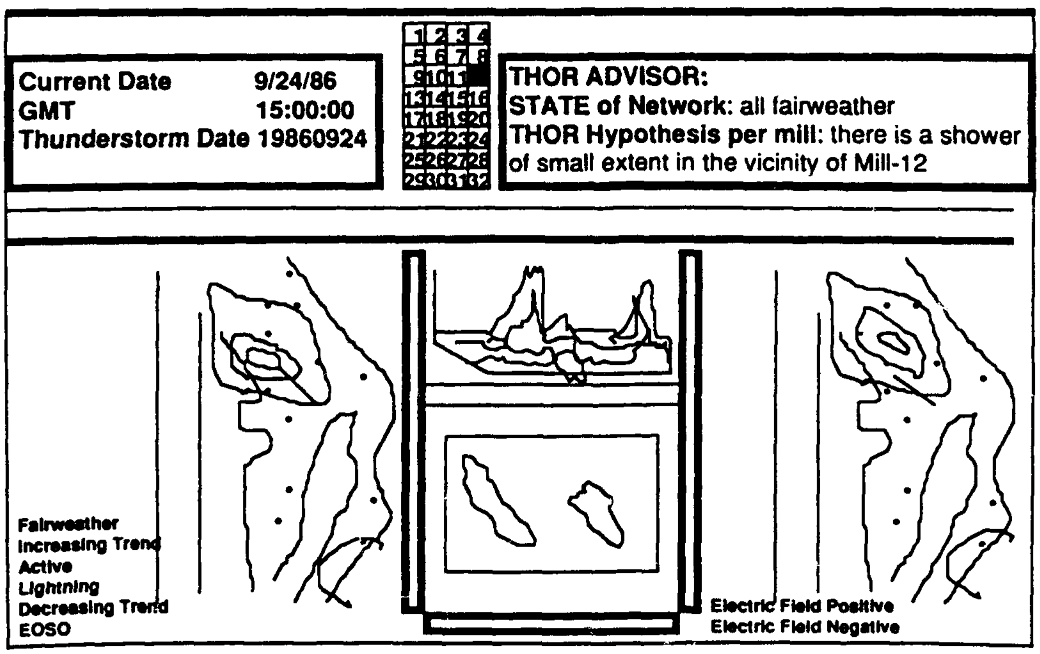

Figure 8. Suggested User Interface Design

In the operational environment, the real time interface of THOR would be the permanent display. If a forecaster or researcher wanted to pursue some post storm analysis, or to use the system to interactively analyze data that is not resident in the system (using the **Interactive** option of the Data Advisory menu as shown in Figure 7) that user would need access to another version of the system so as to not interrupt the real time version.

The details for an operational version of the system are beyond the scope of this report. However, the current version of the system is complete with an on-line tutor (see Figure 7), off-line system documentation [7, 8, 9, 10, 11, 12, 13], and two videotapes (a promotional videotape and a training videotape). Using this information, future potential operators, developers, and system maintainers should be able to comprehend how the system was built, how to operate it, and how to maintain and extend it.

6.1 INTERACTIVE USAGE OF THE DATA ADVISOR

In May 1993, project personnel received information that two of the field mills showed some activity during a recent launch and that there was some LDAR phenomenon observed. The environmental situation is described in the following narrative:

"When the shuttle was going through the speed transition from subsonic to supersonic speeds at about 24000 feet altitude, a burst of pulses was picked up by the LDAR. They occurred over a 3-second period. The shuttle was above existing cloud layers in a layer with very low relative humidity and in which there was some directional and speed shear in the wind. These meteorological factors may well not have been relevant. When KSC meteorologists viewed 35mm movies of the launch there was evidence of *blooming* around the vehicle and of the pressure wave that we perceive as *sonic boom* at that point in the ascent. There apparently is some conjecture that pressure reduction might have permitted some kind of corona discharge that caused the LDAR radiation pulses. The field mill data showed no field excursions that field mill analysts could associate with the LDAR pulses. However, Field Mill 9, which is near the beach at a site in between Pads 39A and 39B did show a trend away from the fairweather values after the launch that the field mill analysts attributed to the presence of the exhaust plume in the vicinity. Field Mill 7, which is located on SR402 north-northwest of the shuttle runway had a number of curious effects, beginning even before the launch. There was a 60-cycle oscillation in the trace, whose amplitude was fairly constant, and there were some other departures from the previous fairweather values. This mill is near a camera pad. The field mill analysts attributed these effects to a generator having been run in the vicinity and to man-induced phenomena, and concluded that nothing showed up in the field mill traces that were related to the LDAR pulses. That is not too surprising since the shuttle was about 8 km high and 10 km down trajectory from the launch pads, and from any field mills, at the time. It would have been unlikely that the field mills would have picked up a small signal because of the distance to the shuttle. There were no anomalies observed in the Shuttle behavior or in any of the elements measured that have to do with vehicle health and functioning."

Provided with some electric field mill data from EG&Gs DPS (Data Processing System), the Data Advisor was run and returned rather interesting hypothesis. For example, given the following: mill-7 was recently calibrated, mill-7 is a boundary mill, mill-7 deviated from fairweather, slope of deviation was not greater than 50 v/m/m, zero was not crossed my mill-7, zero was not crossed by any other mills, and, the direction of deviation from fairweather was an increasing magnitude; the results generated by the Data Advisor are: mill-7 might be bad (belief .45), check the influence of sea breeze (belief .45), possibility of smoke or other field increasing factor such as fog or advection of ions; check neighboring mills (belief is .09). The result that mill-7 might be bad confirms the speculations described in the narrative.

Given that the scenario indicates that mill-9 has also deviated the Data Advisor concludes: check for influence of sea breeze (belief .45), possibility of smoke etc., (belief .18), check calibration of mill-9 (belief .18), there is an anvil overhead, check that all mills that are perturbed from fairweather are of the same polarity (belief .18), AND with the additional information about mill-7 being a **land** mill, the system further concludes: an anvil has expanded over the network from a distant storm, and a new storm is forming over the mill, watch for a change in the slope or a polarity reversal. The hypothesis of "an anvil has expanded overhead" is the most similar explanation the Data Advisor can currently offer given that it contains no *knowledge* about the effect of an exhaust plume from a shuttle launch on a field mill.

The domain specialists were satisfied that the system was able to generate these conclusions.

SECTION 7

CONCLUSION

THOR's Data Advisor generates hypotheses about the storm data as it is received in one minute database chunks from the SIM. The SIM identifies signatures in the data through use of its library of signature detection algorithms. The SIMs signal processing and mathematical capabilities are also employed to generate state vectors, average values, slopes, and graphics displays.

THOR provides multiple ways for a forecaster or researcher to analyze data from archived or current storms. The Data Advisor offers methods for analyzing and generating hypotheses about field mill activity with or without any internal associated data. The Data Advisor also provides a tutor that instructs the user on using and maintaining the system.

The Correlator provided through the Data Advisory module, utilizes the previous conclusions of the Data Advisor in order to monitor the state of the network and the changes in hypotheses as the storm matures and decays. Through the Correlator, the user can query by time, mill, or state (storm phase) to determine how mills are correlated. The Correlator is not a real time tool. Its prime value is as a post analysis tool. In a real time operation, the human eye is better able to correlate mill activity across the network by viewing the contour plots and mill color changes that are displayed on the one-minute updated map plots generated by the SIM.

The long term goal of this work is the realization of a real time system that accurately and automatically detects and identifies signatures associated with thunderstorm atmospheric behavior, and the boundary of these behaviors. Knowledge about the interpretation of multiple storm types needs to be obtained in order to provide a more generalized and robust system and to enable THOR to become part of a more thorough weather forecasting system.

LIST OF REFERENCES

1. Merrem, F. H., and R. H. Brady, "Evaluating an Expert System for Forecasting," Proceedings Fourth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology, Anaheim, CA., American Meteorological Society, 259 - 261, 1988.

2. Stewart, T. R., W. R. Moninger, J. Grassia, R. H. Brady, and F. H. Merrem, "Analysis of Expert Judgment in a Hail Forecasting Experiment," Weather and Forecasting, Vol. 4, March 1989.

3. Peak, J. E., and P. M. Tag, "An Expert System Approach for Prediction of Maritime Visibility Obscuration," Monthly Weather Review, Vol. 117, December 1989.

4. Beller, A. E., and S. P. Stafford, "Space Transportation System Meteorological Expert" Proceedings, 24th Space Congress, Cocoa Beach, Canaveral Council of Technical Society, 8(1) - 8(4), 1987.

5. Quinlan, J. Ross, "Learning Efficient Classification Procedures and Their Application to Chess End Games," Machine Learning: An Artificial Intelligence Approach, ed. by R. Michalski, J. Carbonell, and T. Mitchell, Morgan Kaufmann Publishers, Los Altos, CA., 1983.

6. Thomas, W., and W. Hapgood, "1st Class Reference Manual," 1st Class Expert Systems, Inc., MA, 1989.

7. Maier, L. M., and Capt. T. R. Strange, "Electric Field Mill Training Handbook," January 1988.

8. Mulvehill, A. M., R. A. Marcotte, S. S. El-Sheikh, and E. R. Laderman, "THOR - An Electric Field Mill Data Analysis System, Year End Report 1991 - 1992," MITRE MTR 92B0000033, February 1992, The MITRE Corporation, Bedford MA.

9. Nicholson, J. R., and W. Jafferis, "Atmospheric Sciences Program at NASA Kennedy Space Center," AIAA-88-0197, from AAAI 26th Aerospace Sciences Meeting, January 11 - 14, 1988, Reno, NV., 1988.

10. Nicholson, J. R., L. M. Maier, and J. Weems, "Lightning Threat Extent of a Small Thunderstorm," AIAA-88-1098, from AAAI 26th Aerospace Sciences Meeting, January 11 - 14, 1988, Reno, NV., 1988.

11. Nicholson, J. R., L. M. Maier, and A. M. Mulvehill, "Interpreting Electric Field Mill Data Through Knowledge-Based Techniques," American Geophysics Union, December 9 - 13, 1991, San Francisco, CA., 1991.

12. Nicholson, J. R., and A. M. Mulvehill, "Expert Knowledge Techniques Applied to the Analysis of Electric Field Mill Data," Conference on Atmospheric Electricity, October 22 - 26, 1990, Kananaskis Provincial Park, Alberta Canada, 1990.

13. Rompalo, J. T., "Determination and Representation of Electric Charge Distributions Associated With Adverse Weather Conditions," 1992 Research Reports, NASA/ESEE Summer Faculty Fellowship Program, NASA CR-191004, ed. by Loren Anderson et. al., pp. 391 - 420, September 1992.